Title: Learning What We Don’t Care About: Regularization with Sacrificial Functions

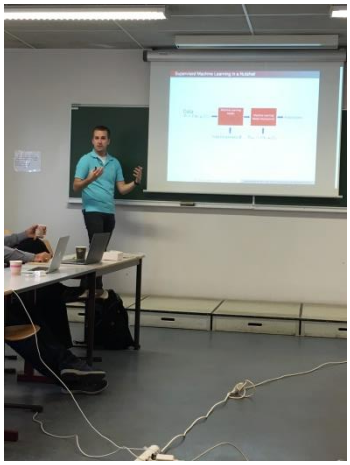

Speaker: Gregory Ditzler, University of Arizona

Abstract: The traditional machine learning paradigm focuses on optimizing an objective. This task of model selection is carried out by adjusting the free parameters of a model that minimizes; however, many times we do not assess the performance of the model using the optimized objective. Furthermore, choosing a poor set of free parameters for the objective function could lead to an unintentional overfitting of the data. Regularization techniques have been shown to mitigate the effects of overfitting; however, these techniques are not formulated very well for meta-, or black box optimization problems. In this talk, I will discuss a how we can use the theory of anti-training for generating sacrificial functions to improve a model's performance. These sacrificial functions incorporate new classifier-independent objectives into the task of choosing a models' free parameters.